

ILLUSTRATED TECHNICAL EDITION

MASTER TECHNOLOGY & ARCHITECTURE COMPARISON

IQAI V3 vs Palantir

Executive assessment, detailed findings and build priorities

The central finding

Palantir is materially ahead in demonstrated operational maturity. IQAI's credible opportunity is a narrower architecture for evidence, qualified cognitive roles, independent challenge and sovereign institutional intelligence.

Architecture is not implementation. The report preserves the distinction between demonstrated capabilities, partial work, architected systems and proposed capabilities.

Recommended direction

CATCH UP → DIFFERENTIATE → SURPASS

Keep the V3 topology. Build reliable foundations, implement the governance contracts, then test whether they produce measurable operational advantages.

SEPTEMBER 2026 · SIX ARCHITECTURE DIAGRAMS

Document assembly and presentation only. The source report's findings, classifications, priorities and references are retained; this edition adds no new technical research or implementation verification.

Read this report

01 Executive assessment	3
02 Detailed architecture comparison	7
03 What to learn and what to integrate	19
04 IQAI differentiation and technology gaps	24
05 V3 changes and Palantir interoperability	32
06 Delivery roadmap and positioning	36
07 External architecture benchmarks	40
08 Sources and reference links	43

The six diagram plates

01 Palantir in One Page	6
02 Palantir — The Operating Loop	9
03 Palantir — What the Ontology Actually Does	11
04 IQAI vs Palantir — Different Centers of Gravity	23
05 IQAI — Intelligence Operating Architecture	31
06 IQAI — Engineering Build Map	35

Reading edition: layout and navigation revised; detailed findings, classifications, priorities and source references retained from the supplied report.

All six diagrams are reproduced from the selected PowerPoint sources; Figures 04–06 use the readable revisions. Full-page landscape plates keep the visual detail separate from the portrait reading pages.

01 Executive assessment

Evidence basis and maturity

This review uses the three supplied IQAI documents as the controlling description of IQAI: IQAI Master Architecture — Adversarial Design Review and Master Architecture V3 Recommendation, IQAI Master Architecture Evidence Addendum, and IQAI V3 Model–Brain Role–Specialist Capability Separation: Adversarial Benchmark.

The important evidentiary distinction in those documents is preserved here: IQAI contains a mixture of built/proven components, partial/advanced components, architected components and proposed implementation. In particular, the Evidence Addendum upgrades Academy from an early hypothesis to a substantially architected subsystem, but explicitly does not establish the persistent Academy, qualified local Brain, qualification platform, autonomous Scout or full multi-brain federation as production-complete.

Palantir is assessed only where current public documentation demonstrates a capability in Foundry, AIP, Gotham or Apollo. A capability that could plausibly be custom-built with Palantir Functions, agents or external code is not treated as equivalent to a documented platform capability. Beta or preview capabilities are identified where material.

Where Palantir leads

The adversarial result is more demanding—and more useful—than a claim that IQAI is simply “a better Palantir.”

Palantir today is the substantially more complete operational technology platform. Foundry has a mature Ontology, broad data connectivity, lineage, object/action semantics, models, functions, applications, APIs and granular authorization. Gotham adds operational/geospatial constructs and Gaia. AIP adds LLMs, agents, tools, model choice, evaluations and increasingly heterogeneous customer-hosted models. Apollo provides a deployment/control plane spanning cloud, on-premise, edge, disconnected and air-gapped environments. Palantir also now supports real-time Ontology peering between separate enrollments and Gotham file peering. [1]

IQAI V3 is therefore not differentiated merely because it has an ontology-like world model, model-independent knowledge, multiple AI models, agents, tools, human approvals, specialists, evaluation, geospatial capabilities, sovereign deployment aspirations or federation. Palantir has credible equivalents to most of those ideas, often in shipping form. Palantir even has a component called the Object Data Funnel, although it performs Ontology indexing/write orchestration rather than the hostile-input preservation and evidentiary intake function assigned to IQAI FUNNEL. [2]

The model-role distinction

The most important correction to an overly optimistic IQAI competitive narrative is around $\text{MODEL} \neq \text{BRAIN ROLE} \neq \text{SPECIALIST CAPABILITY}$. Palantir clearly supports replaceable and heterogeneous models: its current Registered Models architecture can use Palantir-provided models, externally hosted customer models or self-hosted models running on customer compute, including on-premise and air-gapped deployments. Models can be selected in multiple AIP applications; functions can wrap deployed models; multiple model calls can be orchestrated; and AIP Evals can compare alternative models or functions. [3]

So basic model/agent separation is not an IQAI moat. Nor is “institutional knowledge survives a model swap” by itself: Palantir can retain Ontology data, documents, Functions, Actions, application configuration and tool definitions while changing the underlying model. The IQAI distinction survives only in a narrower form: V3 proposes to make role, authority domain, evidence requirements, world-state eligibility, tool permissions, institutional competency and qualification/requalification first-class external contracts, with a model unable to occupy a BrainRole merely because it is technically available. The supplied IQAI model-role sub-study reached the same conclusion.

That narrower architecture is genuinely interesting because public Palantir documentation reviewed here demonstrates model registration, model selection, evaluation and controlled change, but not a common platform-wide invariant equivalent to an IQAI BrainInstance Passport whose eligibility is explicitly conditioned on the entire binding of BrainRole + ModelArtifact + runtime + configuration + knowledge version + tool set + permissions + context policy + specialist access, with expiry/requalification affecting runtime dispatch. Palantir's AIP Evals and AIP Evolve move quite close to the problem—the latter can evaluate proposed system changes such as model substitutions rather than assuming interchangeability—but they do not publicly establish the exact IQAI qualification-and-authority architecture. [4]

Independent challenge and the remaining opportunity

Likewise, Mission Brain + Reviewer/Challenger is not unique merely because it is multi-agent. Current agent architectures can coordinate specialized agents, and Palantir can publish agents as callable Functions/tools and compose agentic systems. The potential IQAI distinction arises only if the Reviewer/Challenger is structurally independent, selectively mandatory, independently qualified, unable to inherit the Mission Brain's authority, and connected to explicit disagreement/arbitration records rather than being another prompt in an application. Palantir's current documentation demonstrates composition; it does not demonstrate that constitutional invariant. [5]

IQAI's other defensible architectural differences are similarly combinations, not isolated primitives: a typed epistemic layer separating artifacts, observations, claims, evidence, accepted baseline assertions, hypotheses and recommendations; explicit baseline/reconstruction/hypothesis/forecast world-state separation; domain-authoritative specialists spanning AI, deterministic engines, scientific models and human SMEs; institutional competency that survives model replacement; qualification as a routing dependency; and a sovereign federation intended to span heterogeneous infrastructure rather than only instances of one vendor platform. These are constitutional goals in V3—not yet proof of production superiority.

The practical conclusion is therefore:

Palantir has already solved much of the platform engineering that IQAI would otherwise spend years rediscovering. IQAI should not try to win by recreating Foundry, Gotham or Apollo. It should use established infrastructure patterns and commodity/open components underneath a much thinner IQAI-owned constitutional layer: evidence semantics, world-state semantics, cognitive roles, specialist authority, qualification, controlled action, institutional competency and sovereign federation.

Five executive questions

What is it strongest at today?

PALANTIR An integrated, operational enterprise/mission platform joining data, Ontology, applications, AI, security and deployment. [6]

IQAI V3 An assurance-oriented architecture for governed world state, evidence, specialist intelligence, replaceable cognitive roles and controlled human authority. (IQAI V3 Review, “Reconstructed master architecture” and “Foundational primitives”.)

Who is technologically ahead today?

PALANTIR Palantir, materially, across implementation maturity, connectors, Ontology, security, developer tooling, deployment, operations, observability, federation and geospatial integration. [7]

IQAI V3 Significant architecture exists, but several foundations remain partial or P0/P1 implementation work. (IQAI V3 Review, gap and build-priority sections; Evidence Addendum.)

Is IQAI's multi-model architecture unique?

PALANTIR No. Palantir supports multiple providers, customer/self-hosted models, tool-capable models, model switching and evaluation. [8]

IQAI V3 The potentially distinct element is qualified role occupancy and specialist authority, not model agnosticism itself. (IQAI Model–Brain Role benchmark.)

Where can IQAI be genuinely different?

PALANTIR Palantir generally organizes capabilities around Ontology, Functions, Actions, models, agents and governed applications. [9]

IQAI V3 Epistemic classes + explicit alternative worlds + qualified cognitive roles + independently governed challengers + specialist authority + portable institutional competency.

Where could IQAI eventually surpass?

PALANTIR Palantir remains the stronger baseline platform today.


IQAI V3 In assurance-critical, sovereign intelligence environments if qualification, evidence-grounded replay, model-independent competency and cross-vendor cognitive federation become real runtime invariants rather than documentation.

FIGURE 01 / PALANTIR PUBLIC ARCHITECTURE

PALANTIR IN ONE PAGE

01 / ARCHITECTURE

High-level operating architecture based on publicly available information



Conceptual operating model, not exact runtime topology. Foundry includes the integration and Ontology foundation as well as applications. AIP is optional; Apollo denotes software delivery and management.

Source: Palantir-In-One-Page.pptx | Supplied diagram and interpretation notes, 05 September 2026.

02 Detailed architecture comparison

24 areas • each finding retains the source classification

How to read the classifications

The classification below is deliberately strict.

PALANTIR STRONGER means both sides have a meaningful concept but Palantir has the more complete demonstrated implementation.

IQAI GAP means V3 itself identifies missing or incomplete implementation serious enough to prevent operational parity.

IQAI STRUCTURALLY DIFFERENT does not mean “better.” POSSIBLE IQAI ADVANTAGE means the distinction could become competitively meaningful only after implementation and validation.

Overall conclusion

The adversarial review removes several easy differentiation claims. Palantir has already crossed the lines from data platform → governed world model → operational action → AI → deployment/federation. It has also crossed much further into heterogeneous models and agent composition than a comparison based on older Palantir descriptions would suggest. [33]

What remains for IQAI is not “Palantir without vendor lock-in” and not “Palantir plus multiple LLMs.” The defensible architectural thesis is a constitutional intelligence layer in which epistemic status, world state, cognitive role, domain authority, institutional competency, qualification and human authorization remain externally governed and portable across changing models, specialist implementations and infrastructure providers.

Comparison index


Architecture area	Source classification
Universal intake / multimodal ingestion / FUNNEL	PALANTIR STRONGER
Data fabric and connectors	IQAI GAP
Evidence, provenance and lineage	IQAI STRUCTURALLY DIFFERENT
Digital World vs Palantir Ontology	PALANTIR STRONGER
Entity, relationship and event modelling	PALANTIR STRONGER
GIS / Spatial / GEOINT / terrain / imagery / 3D	PALANTIR STRONGER
Historical reconstruction / time / alternative worlds / simulation	IQAI STRUCTURALLY DIFFERENT
Sensors, cameras, streaming and edge data	IQAI GAP
TRACE / investigation and analytical workflows	PALANTIR STRONGER
AB / cognitive federation	ROUGHLY EQUIVALENT
Mission / Reviewer / Challenger intelligence	POSSIBLE IQAI ADVANTAGE
Specialist AI, deterministic engines, scientific models and human SMEs	IQAI STRUCTURALLY DIFFERENT
MODEL vs BRAIN ROLE vs SPECIALIST CAPABILITY	POSSIBLE IQAI ADVANTAGE
Institutional knowledge and memory	ROUGHLY EQUIVALENT
Academy / Scout / qualification / requalification	POSSIBLE IQAI ADVANTAGE
Human authority / recommendations / authorized actions	PALANTIR STRONGER
Workflow durability and orchestration	IQAI GAP
Security, identity, permissions and zero trust	PALANTIR STRONGER
Sovereign / local / disconnected deployment	IQAI GAP
Federation across organizations	IQAI GAP
Deployment, updates, DevOps and rollback	IQAI GAP
Observability, audit and replay	IQAI GAP
APIs / SDKs / developer ecosystem	IQAI GAP
Scalability, reliability and enterprise maturity	IQAI GAP

FIGURE 02 / PALANTIR PUBLIC ARCHITECTURE

PALANTIR — THE OPERATING LOOP

02 / OPERATING MODEL

How data becomes governed operational context, analysis, workflow and action



01 Universal intake / multimodal ingestion / FUNNEL

PALANTIR STRONGER

PALANTIR Data Connection supports many enterprise/streaming sources; Foundry supports media such as images/audio/documents; streaming pipelines have dedicated primitives. Palantir's Ontology backend also has an “Object Data Funnel,” but that is an indexing/write service, not equivalent to IQAI's evidentiary intake gateway. [10]

IQAI V3 FUNNEL receives, quarantines, preserves, detects, decomposes, normalizes and packages material while retaining immutable originals; high-rate sources are deliberately separated into a Sensor/Stream Gateway. (V3 Review, “FUNNEL must become the formal evidence boundary”).)

FINDING Palantir wins breadth and operational maturity. IQAI's potentially useful distinction is preservation/epistemic intake, not generic ingestion. Do not market FUNNEL as a connector advantage until built at scale.

02 Data fabric and connectors

IQAI GAP

PALANTIR Broad connector catalogue spans databases, cloud stores, ERP/CRM, Kafka/Kinesis, messaging, REST, IoT/OPC-UA and geospatial systems; platform supports batch, streaming and other integration patterns. [11]

IQAI V3 V3 assumes universal references and intake/provider federation but does not evidence a Palantir-scale connector estate.

FINDING This is commodity-but-critical plumbing. IQAI should integrate existing connector technology, not make connector count a core R&D programme.

03 Evidence, provenance and lineage


IQAI STRUCTURALLY DIFFERENT

PALANTIR Foundry tracks data lineage and change history and provides comprehensive audit logging. Ontology Actions create controlled edits and historical operational records. [12]

IQAI V3 V3 explicitly separates Artifact, Derived Artifact, Observation/Report, Claim, Evidence Link, Derived Result, Accepted Baseline Assertion, Hypothesis/Reconstruction/Simulation/Forecast, Recommendation, Decision and Action/Outcome. (V3 Review, epistemic model.)

FINDING Palantir is far stronger operationally in lineage/audit. IQAI's distinction is that provenance is intended to become epistemic and evidentiary, not merely data/process lineage. That difference is valuable only if every Brain and Specialist actually emits those contracts.

FIGURE 03 / PALANTIR PUBLIC ARCHITECTURE



04 Digital World vs Palantir Ontology

PALANTIR STRONGER

PALANTIR The Ontology is an operational/digital-twin layer connecting datasets/models to real-world objects, properties and links, plus “kinetic” Actions, Functions and dynamic security. It has a multi-service backend for indexing, querying and controlled writes. [13]

IQAI V3 Digital World is a governed logical baseline of canonical identities, relationships, accepted state, geometry and time; V3 explicitly says it should be logically unified but physically polyglot.

FINDING The Ontology is the clearest Palantir benchmark IQAI must respect. “We have a Digital World” is not differentiation. IQAI can differentiate only through evidence/world-state semantics and cross-vendor federation.

05 Entity, relationship and event modelling

PALANTIR STRONGER

PALANTIR Ontology object/link types, interfaces, Actions and Functions are first-class and strongly integrated into operational applications; Gotham integration also maps entity/event/document categories. [14]

IQAI V3 Canonical identity, relations and event/mission fabric are foundational primitives, but durable event architecture and distributed consistency remain P0 work.

FINDING IQAI should borrow the semantic + kinetic lesson: a world model without governed behaviour attached becomes just another graph/database.

06 GIS / Spatial / GEOINT / terrain / imagery / 3D

PALANTIR STRONGER

PALANTIR Foundry/Gotham integrate Ontology objects with Gaia, geotemporal series and standard WFS access; Foundry has geospatial transformations and media support. [15]

IQAI V3 V3 separates foundational spatiotemporal semantics from a specialist Spatial/GEOINT capability; supplied maturity evidence marks portions of spatial, ground, 3D and imagery federation as partial.

FINDING Palantir is ahead; Esri is an even stronger specialist benchmark for the deepest GIS/imagery/3D portions. IQAI should preserve vendor-neutral spatial semantics rather than attempt to out-build mature GIS engines.

07 Historical reconstruction / time / alternative worlds / simulation

IQAI STRUCTURALLY DIFFERENT

PALANTIR Foundry supports branching, beta Ontology Scenarios for isolated what-if changes and Vertex scenario/model workflows; these give Palantir real scenario and simulation capability. [16]

IQAI V3 V3 makes BASELINE, RECONSTRUCTION, HYPOTHESIS, SIMULATION and FORECAST/SCENARIO explicit epistemic/world-state modes with contamination boundaries.

FINDING Palantir decisively disproves any claim that it lacks alternate scenarios. IQAI's remaining difference is epistemic typing and evidence linkage between worlds, especially historical reconstruction versus speculative world state.

08 Sensors, cameras, streaming and edge data

IQAI GAP

PALANTIR Foundry Streams supports low-latency ingestion, persistence and restart/checkpoint patterns; Gotham has observation/geotemporal constructs; Palantir's broader platform supports edge/disconnected deployment. [17]

IQAI V3 V3 correctly adds a separate Sensor/Stream Gateway so FUNNEL does not receive every video frame or telemetry point. Several existing IQAI sensor/camera components remain incompletely federated.

FINDING The architectural correction is right; the operational benchmark should be Anduril Lattice as much as Palantir.

09 TRACE / investigation and analytical workflows

PALANTIR STRONGER

PALANTIR Foundry Object Explorer searches, filters, maps and analyzes Ontology object sets; Gotham brings operational/investigative capabilities and shared Foundry/Gotham representations. [18]

IQAI V3 TRACE has a proven investigation foundation in the supplied evidence, while full hypothesis/reconstruction federation remains partial.

FINDING IQAI should not assume investigation workflow is unique. Its opportunity is to bind investigation conclusions more tightly to explicit claims, evidence, alternative-world state and reviewer challenge.

10 AB / cognitive federation

ROUGHLY EQUIVALENT

PALANTIR Palantir agents can be exposed as callable capabilities and composed with other agents/tools; pro-code support allows external agent SDKs in Foundry. [5]

IQAI V3 V3 turns AB from a “Central Brain” into governed cognitive federation with Mission Brain, Challenger/Reviewer, optional roles, deterministic verification and routing policy.

FINDING The generic idea—multiple specialized AI participants—is no longer differentiating. IQAI must compete on governance of federation, not existence of federation.

11 Mission / Reviewer / Challenger intelligence

POSSIBLE IQAI ADVANTAGE

PALANTIR Palantir can technically construct writer/reviewer, orchestration and agent-as-tool patterns, and AIP Evals can independently test outputs. [19]

IQAI V3 Reviewer/Challenger is intended as a separate role with independence, separate authority ceiling, qualification, disagreement records and arbitration.

FINDING No reviewed Palantir public documentation established this as a platform-wide mandatory independence contract. But it is not an advantage until IQAI proves independence, routing policy and measurable error reduction in production.

12 Specialist AI, deterministic engines, scientific models and human SMEs

IQAI STRUCTURALLY DIFFERENT

PALANTIR Foundry Functions can execute arbitrary business logic; model functions integrate deployed models with Ontology operations; Modeling Objectives operationalize models; external systems can be invoked. [20]

IQAI V3 SpecialistCapability is explicitly domain-authoritative and may be AI, deterministic code, scientific modelling or human expertise, with authority/evidence/limitations declared independently of implementation.

FINDING Palantir has the components. IQAI's potential distinction is one uniform semantic contract covering authority and evidentiary fitness across very different intelligence types.

13 MODEL vs BRAIN ROLE vs SPECIALIST CAPABILITY

POSSIBLE IQAI ADVANTAGE

PALANTIR Registered Models abstract multiple model providers/sources; models can be changed without rebuilding the whole application; model functions provide stable callable boundaries; Evals compare models/functions. [21]

IQAI V3 ModelArtifact ≠ BrainRole ≠ BrainInstance ≠ SpecialistCapability ≠ Tool ≠ Knowledge ≠ Memory; BrainRole survives model replacement and BrainInstance is the qualified binding. (Model–Brain Role benchmark; V3 Review.)

FINDING Model-role separation itself is roughly equivalent and not unique. The possible advantage is the additional binding of authority + knowledge version + tools + permissions + world-state scope + qualification and making that binding govern dispatch eligibility.

14 Institutional knowledge and memory

ROUGHLY EQUIVALENT

PALANTIR Ontology, documents, object state, Functions, tools and chatbot/agent state live outside foundation-model weights, permitting application context to survive model replacement. [22]

IQAI V3 IQAI explicitly separates Knowledge, procedural memory, operational/episodic memory and model context; the Academy retains sources, structured knowledge, methods, assessments and qualification history.

FINDING “Knowledge survives a model swap” is not enough to distinguish IQAI. The stronger IQAI thesis is a portable institutional competency package, including tests and authorization criteria—not persistence alone.

15 Academy / Scout / qualification / requalification

POSSIBLE IQAI ADVANTAGE

PALANTIR AIP Evals supports test suites, expected outputs, custom/built-in evaluators, multi-target comparison and repeated runs; AIP Evolve can test proposed system changes including model substitutions before merging them. [23]

IQAI V3 Academy owns source qualification, structured knowledge, curriculum, assessments, competency, qualification/requalification and knowledge gaps; affected changes can force REQUALIFICATION_REQUIRED; independent TEVV remains separate. (Evidence Addendum.)

FINDING This survives adversarial testing only in its strongest form: qualification is an expiring runtime entitlement to fill a role/domain, not merely an evaluation result. Academy is substantially architected but still mostly proposed operationally.

16 Human authority / recommendations / authorized actions

PALANTIR STRONGER

PALANTIR Ontology Actions are governed edits; AIP Chatbot tools can require confirmation; AI FDE documentation imposes explicit approval for sensitive/mutating operations. [24]

IQAI V3 V3's ActionEnvelope would bind proposer, evidence, risk class, authorization roles/quorum, policy result, signatures, expiry, execution adapter, idempotency, outcome and post-action verification.

FINDING Palantir has the stronger working system today. IQAI's ActionEnvelope could become a stronger assurance primitive for consequential operations, but it must first be implemented end-to-end.

17 Workflow durability and orchestration

IQAI GAP

PALANTIR Palantir Automate has event/schedule/object triggers, retries, fallbacks and execution/observability support. [25]

IQAI V3 V3 explicitly identifies durable mission state, retries, deadlines, cancellation, human waits, idempotency and replay as missing foundational infrastructure.

FINDING Do not build an ad-hoc workflow engine. Temporal is the stronger architectural benchmark for durable execution.

18 Security, identity, permissions and zero trust

PALANTIR STRONGER

PALANTIR Palantir documents granular Ontology permissioning, row/property-level controls, markings/mandatory controls, permission propagation, user-scoped AI access and a zero-trust service architecture. [26]

IQAI V3 V3 requires human/workload/node identity, ABAC, secrets isolation, sandboxed tools, role-scoped AI permissions, classification/jurisdiction segregation, signed supply chain and tamper-evident audit.

FINDING This is a major Palantir strength and should be treated as a minimum bar, not an area where architectural prose creates parity.

19 Sovereign / local / disconnected deployment

IQAI GAP

PALANTIR Apollo manages software across heterogeneous environments including private infrastructure and disconnected/air-gapped locations; AIP can self-host customer models on-premise or air-gapped. [27]

IQAI V3 Sovereign Node & Federation is a V3 primitive, including local operation, independent control/data planes, signed updates and disconnected reconciliation, but the trust/disconnected architecture remains incomplete.

FINDING Architecturally aligned, operationally far behind. Apollo is one of the most important systems IQAI should study.

20 Federation across organizations

IQAI GAP

PALANTIR Peer Manager synchronizes Ontology objects/links between separate enrollments, supports mappings and directionality, can peer Gotham files, and now has reusable peer profiles with security constraints. [28]

IQAI V3 IQAI intends sovereign-node federation across clients/jurisdictions, with evidence/world state/capabilities synchronized under policy; protocols, reconciliation and disconnected trust remain incomplete.

FINDING Palantir's peering materially raises the competitive bar. IQAI differentiation requires heterogeneous/cross-vendor federation, not merely multiple IQAI nodes.

21 Deployment, updates, DevOps and rollback

IQAI GAP

PALANTIR Apollo provides centrally governed deployment across independent environments, promotion constraints, maintenance windows and rollback/roll-off mechanisms. [29]

IQAI V3 V3 requires signed, staged, rollback-capable and offline updates, but operational implementation is not evidenced at Apollo maturity.

FINDING Do not build a proprietary “mini-Apollo” unless required by a sovereign niche. Use OCI/Kubernetes/GitOps-style interfaces and make deployment backends replaceable.

22 Observability, audit and replay

IQAI GAP

PALANTIR Foundry exposes distributed traces, logs, metrics and run history across models, functions, Actions and automations; audit logs record who/what/when/where for platform actions. [30]

IQAI V3 V3 requires a mission correlation identity spanning intake → evidence → world reads → brains → specialists → disagreement → authorization → action → outcome, plus forensic replay.

FINDING Palantir has stronger operational observability. IQAI should add domain semantics above OpenTelemetry rather than invent a telemetry stack.

23 APIs / SDKs / developer ecosystem

IQAI GAP

PALANTIR REST APIs, platform SDKs and generated Ontology SDKs allow applications to manipulate platform and Ontology resources; OSDK generation supports Python, Java and TypeScript in the dev toolchain. [31]

IQAI V3 SDK/contract harness is P1 in V3; capability manifests and envelopes are intended to make extension systematic.

FINDING A contract is not an ecosystem. IQAI needs a genuinely low-friction SDK, schema generation, local test harness and conformance suite.

24 Scalability, reliability and enterprise maturity

IQAI GAP

PALANTIR Foundry/AIP/Gotham/Apollo are integrated shipping platforms with specialized backend services, observability, data/AI tooling and deployment management across varied environments. [32]

IQAI V3 V3 explicitly identifies registry centralization, provenance growth, high-rate streams, synchronous FUNNEL processing, model fan-out, disconnected consistency and DR as scaling risks requiring redesign.

FINDING This is the largest aggregate difference. IQAI does not need Palantir's total breadth, but it needs evidence that its chosen slice survives production load, failure, upgrade, compromise and disconnection.

03 What to learn and what to integrate

Where Palantir is materially ahead

Semantics tied to operations

The largest Palantir advantage is integration of semantics with operations. The Ontology is not merely an entity graph: objects and links are coupled to Actions, Functions, dynamic security and end-user workflows. That is a major architectural lesson for IQAI's Digital World. IQAI should avoid creating a pristine canonical world representation that every useful operational capability then circumvents. Governed world objects must be addressable by computation, policy and action using stable contracts. [\[34\]](#)

Production infrastructure

The next advantage is production plumbing. Palantir has broad Data Connection support, batch/streaming primitives, model infrastructure, developer SDKs, audit, distributed observability and an extensive permissions model. None is glamorous, but together they determine whether an architecture becomes a usable platform. [\[35\]](#)

Security across the platform

Palantir is particularly far ahead in security as a platform invariant. Its Ontology permissions, mandatory controls/markings, propagation model and AI execution under user-scoped permission boundaries demonstrate that security is not simply a perimeter around the AI layer. IQAI's proposed ABAC, workload identity and role-scoped tool access are directionally correct, but they require an equivalent degree of per-resource enforcement. [\[36\]](#)

Sovereign software lifecycle

Apollo establishes an equally important benchmark for software lifecycle under sovereignty constraints: an AI/mission platform has not solved sovereign deployment merely because its containers can run offline. It needs controlled releases, environment policy, health monitoring, promotion, rollback and a way to update many isolated systems safely. [\[37\]](#)

Governed organizational peering

Peer Manager also weakens any IQAI claim that organizational federation is absent from Palantir. Palantir now publicly documents Ontology peering across distinct enrollments, object/property mappings, source/action directionality, file peering and reusable secured peer profiles. IQAI's federation therefore needs to target the harder problem: policy-preserving federation among sovereign nodes that do not necessarily run the same vendor stack, schema version or connectivity regime. [\[38\]](#)

Model lifecycle governance

Finally, AIP's model infrastructure is now a serious direct benchmark. Registered Models can be external or self-hosted and become first-class choices in supported AIP applications; customer infrastructure can even host a proxy/federation layer. AIP Evals makes model/function comparison explicit. This means IQAI must stop treating "replaceable models" as a differentiation statement and turn its Brain Passport into demonstrably stronger lifecycle governance. [\[8\]](#)

Lessons IQAI should independently adapt

Ontology as semantic + kinetic system

Keep Digital World identity/state independent of applications, but attach typed queries, governed Functions/Specialists and ActionEnvelope operations directly to canonical references. [39]

Granular authorization propagates through the stack

Treat tenant, mission, jurisdiction, classification, purpose and world-state labels as non-optional fields propagated through evidence retrieval, Brain calls, Specialist calls and actions. Palantir's fine-grained Ontology controls are the minimum competitive benchmark. [40]

Model change must be operationalized, not philosophized

Give ModelArtifact immutable/versioned identity, capabilities, runtime digest, licence and lifecycle metadata; make model swap a controlled binding change and execute the relevant qualification suites. [41]

Test changes against representative cases

Use Academy assessments + independent TEVV as executable release gates, preserving historical baselines and variance rather than issuing a one-time “qualified” badge. Palantir AIP Evals/Evolve provide strong product patterns to study. [42]

Federation requires mapping and governance, not just transport

IQAI sovereign peering should explicitly map schema/version, authority, evidence labels, action rights, synchronization direction and conflict/reconciliation policy. [43]

Deployment is part of architecture

Define ReleaseManifest, compatibility ranges, migration semantics, signed artifacts, health checks, staged promotion and rollback as constitutional contracts. Apollo is the reference pattern. [44]

Developer experience is a platform feature

Generate typed SDK bindings from IQAI schemas/manifests, supply contract test kits, emulators and reference adapters. Palantir's OSDK is a particularly good conceptual benchmark. [45]

Observability must cross AI and conventional software

One mission trace should connect model invocations, tool calls, evidence references, Specialist versions, policy results and action receipts, built over standard distributed tracing. Palantir already traces models/functions/actions/automations. [46]

What to integrate instead of rebuilding

IQAI should be ruthless about the difference between architectural sovereignty and writing commodity infrastructure itself.

Giant proprietary connector catalogue

BETTER STRATEGY Standard adapters, existing integration frameworks, Kafka-compatible streaming, JDBC/REST/filesystem/CDC conventions

IQAI MUST OWN IntakeEnvelope, EvidenceEnvelope, security labels and source provenance

Foundation-model inference servers

BETTER STRATEGY Support vLLM/Ollama-equivalent or managed endpoints through adapters; even Palantir's current BYOM design recognizes customer-hosted model servers as interchangeable sources. [47]

IQAI MUST OWN ModelArtifact identity, Brain binding, qualification and routing policy

General-purpose workflow engine

BETTER STRATEGY Use a proven durable execution engine or equivalent

IQAI MUST OWN Mission semantics, human-authority steps, ActionEnvelope and IQAI event types

GIS projection/rendering/tiling engine

BETTER STRATEGY Esri/open GIS/OGC-compatible services

IQAI MUST OWN GeometryRef, spatial authority, mission context, evidence links and provider federation

Generic model registry

BETTER STRATEGY Adopt MLflow-compatible concepts/artifact stores where useful

IQAI MUST OWN BrainRole, BrainInstance, Specialist Passport and qualification eligibility

Kubernetes/container scheduler

BETTER STRATEGY Package IQAI as portable signed workloads

IQAI MUST OWN Sovereign-node policy, release manifests, compatibility and offline recovery

Generic IAM/KMS/secrets stack

BETTER STRATEGY Integrate enterprise identity, PKI, HSM/KMS and workload-identity technology

IQAI MUST OWN IQAI authorization semantics and policy enforcement points

Proprietary telemetry transport

BETTER STRATEGY OpenTelemetry-compatible traces/logs/metrics

IQAI MUST OWN Mission/evidence/Brain/Specialist/qualification/action semantic fields

Data warehouse/lakehouse

BETTER STRATEGY Use customer-selected relational, object, graph, search, time-series and lakehouse stores

IQAI MUST OWN Logical Digital World identity, authority and consistency contracts

Another generic multi-agent SDK

BETTER STRATEGY Use modular agent runtimes where appropriate

IQAI MUST OWN IQAI cognitive constitution: roles, independence, budgets, disagreement, qualification and authority


This is one of the most important strategic recommendations in the report. IQAI should own the semantics that make intelligence governable; it should not own every engine on which those semantics run.

FIGURE 04 / IQAI ARCHITECTURE + BUILD DIRECTION

IQAI vs PALANTIR — DIFFERENT CENTERS OF GRAVITY

Two architectures for turning data and intelligence into governed operational action

04 / ARCHITECTURE COMPARISON



Palantir also supports AI governance, human oversight, model choice, permissions and sovereign/disconnected deployment. The comparison shows architectural emphasis, not exclusive capabilities.

Architectural emphasis, not exclusive capabilities or current platform parity. Qualification governs eligibility before dispatch; its downstream visual placement is a simplification, not permission for unqualified reasoning.

Source: IQAI-vs-Palantir-Centers-of-Gravity-Readable.pptx | Supplied diagram and interpretation notes, 05 September 2026.

04 IQAI differentiation and technology gaps

Six surviving architectural distinctions

The strongest surviving IQAI proposition is not any single component. It is the composition of six governance ideas.

Epistemic computing

The first is epistemic computing. IQAI V3 intends to know whether a piece of information is a preserved artifact, an observation, an asserted claim, supporting evidence, a derived result, an accepted baseline assertion, a reconstruction, a hypothesis, a forecast, a recommendation, a decision or an observed outcome. Foundry has lineage, audit and object state, but the IQAI architecture makes these categories part of the intelligence contract itself. That could materially improve forensic review and reduce the common AI failure of collapsing “retrieved,” “inferred,” “assumed” and “accepted” into one undifferentiated context. Palantir’s lineage/action controls remain a strong implementation benchmark. [\[48\]](#)

Qualified cognitive authority

The second is qualified cognitive authority. A foundation model is not “the Mission Brain.” A Mission Brain is a governed job; one qualified BrainInstance happens to fulfil it. A different model may fulfil Reviewer, while a multimodal model performs Perception. Tools, knowledge, memory and permissions remain external. Palantir demonstrably supports the component separation needed for this pattern, so the differentiation begins only where IQAI makes qualification and authority portable contracts rather than application configuration. [\[41\]](#)

Independent challenge

The third is independent challenge as constitutional architecture. A Reviewer that merely receives the Mission Brain’s answer and says “looks good” is not differentiated. A real IQAI challenger must independently reconstruct enough of the evidence path to detect unsupported claims, alternative explanations, domain violations and confidence failures; its disagreements must be explicit records and consequential disagreements must invoke deterministic/human arbitration policy. Multi-agent orchestration itself is common—Microsoft’s current Agent Framework, for example, supplies sequential, concurrent, handoff, group-chat and manager-coordinated multi-agent patterns, including human approval. [\[49\]](#)

Specialist authority

The fourth is Specialist authority distinct from generic intelligence. An atmospheric dispersion model, hazardous-material reference engine, computer-vision detector, geospatial computation, deterministic solver and qualified human SME should not all be treated as interchangeable “tools.” Their manifests can declare what kind of conclusion they are authoritative for, what evidence they require, what conditions invalidate them and what world states they support. Palantir has rich model/function/tool composition; IQAI can differentiate by making domain authority and evidentiary validity portable metadata across all implementation types. [50]

Institutional competency ownership

The fifth is institutional competency ownership. The useful Academy claim is not that knowledge lives outside an LLM—that is already common. It is that the institution owns the source library + structured knowledge + doctrine/procedures + deterministic tools + assessment bank + failure cases + qualification criteria + qualification history, and a replacement model must prove that it can operate that competency package before acquiring authority. Databricks/MLflow already demonstrates how strongly versioned model lineage and model lifecycle can be externalized; the IQAI opportunity is to extend that idea from “which model version?” to “which qualified cognitive binding may perform this institutional job?” [51]

Governed multi-world intelligence

The sixth is governed multi-world intelligence. Palantir has real scenario, branch and simulation mechanisms, so alternative worlds are not unique. IQAI's possible distinction is making every inference explicitly scoped to BASELINE vs RECONSTRUCTION vs HYPOTHESIS vs SIMULATION/FORECAST and preventing speculative state from silently becoming accepted baseline. [52]

None of these claims should be expressed as “Palantir cannot do this.” A sufficiently capable Palantir team could construct many of them with agents, Functions, Ontology objects and evaluation logic. The defensible claim is more precise: IQAI proposes these as universal platform contracts and invariants rather than leaving each application team to construct them. Whether that becomes an advantage depends entirely on implementation quality.

Technology gaps ranked by urgency

P0 → P1 → P2

Priority labels and acceptance criteria are reproduced from the source report. Each entry states why the gap matters and what would count as credible implementation.

P0 / Foundational implementation

Universal reference + epistemic envelope implementation

WHY IT MATTERS Everything else depends on stable references and unambiguous evidence/world-state semantics.

CREDIBLE ENOUGH WHEN Every artifact, claim, observation, derived result and world assertion has stable ID, provenance, producer/version, time, security and world-state fields.

Digital World authority/write/reconciliation contract

WHY IT MATTERS Without it, “governed world” can become an incoherent collection of stores.

CREDIBLE ENOUGH WHEN Logical canonical identity over multiple physical stores; explicit authority owner, concurrency, valid/event/ingest time and reconciliation semantics.

Event + durable Mission Fabric

WHY IT MATTERS LLM context cannot hold operational state.

CREDIBLE ENOUGH WHEN Crash/restart/retry/human-wait/cancel/idempotency/replay demonstrated in long-running missions.

Capability/Brain/Runtime Registry

WHY IT MATTERS Federation cannot be governed through hard-coded service discovery.

CREDIBLE ENOUGH WHEN Signed versioned manifests, local cache, health, location, permissions, authority, qualification state and compatibility.

BrainInstance / Specialist qualification gating

WHY IT MATTERS This is central to the most credible IQAI differentiation.

CREDIBLE ENOUGH WHEN Dispatch requires available $\wedge$ authorized $\wedge$ qualified_for(role, domain, mission class, world state); expiry or relevant change removes eligibility.

Independent Challenger + arbitration runtime

WHY IT MATTERS Otherwise “Reviewer Brain” is just a diagram.

CREDIBLE ENOUGH WHEN Separate execution/context policy, disagreement objects, independent evidence checking, escalation and measured detection benefit.

Security constitution enforced in runtime

WHY IT MATTERS A mission platform cannot retrofit trust later.

CREDIBLE ENOUGH WHEN Human + workload identity, ABAC, least privilege, tainted input isolation, secrets isolation, cryptographic artifact identity and tamper-evident audit.

ActionEnvelope state machine

WHY IT MATTERS The analysis/action boundary is where AI risk becomes physical/organizational consequence.

CREDIBLE ENOUGH WHEN Recommendation → policy → approvals → execution → receipt → observed outcome, with idempotency and revocation/expiry.

Sovereign-node baseline

WHY IT MATTERS Required to make local/disconnected claims credible.

CREDIBLE ENOUGH WHEN Fully local operation for a bounded mission, signed releases, local identity/policy, backup/recovery and controlled reconnect.

End-to-end mission observability

WHY IT MATTERS Without it IQAI cannot prove any of the above.

CREDIBLE ENOUGH WHEN One trace reconstructs source → evidence → state → models → Specialists → challenge → decision → action → outcome.

P1 / Follow-on implementation

Academy persistence, assessment bank and qualification store

WHY IT MATTERS Converts architecture into institutional competency.

CREDIBLE ENOUGH WHEN Versioned source/knowledge/assessment pipeline with reproducible qualification evidence.

Sensor/Stream Gateway

WHY IT MATTERS Required for video/telemetry scale.

CREDIBLE ENOUGH WHEN Hot/cold path, segmentation/windowing, durable offsets/checkpoints and evidence promotion.

Spatial/GEOINT provider federation

WHY IT MATTERS Avoids spatial subsystem fragmentation.

CREDIBLE ENOUGH WHEN OGC/STAC-compatible providers behind common geometry/time/reference contracts.

Cross-node federation protocol

WHY IT MATTERS Needed for sovereign network effects.

CREDIBLE ENOUGH WHEN Trust-domain mapping, schema negotiation, minimization, security labels, conflict metadata and disconnected reconciliation.

Model/runtime lifecycle registry

WHY IT MATTERS Makes replacement operationally safe.

CREDIBLE ENOUGH WHEN Immutable digests, licences, runtime compatibility, eval history, rollback and expiry.

Persistent memory architecture

WHY IT MATTERS Necessary for continuity without model lock-in.

CREDIBLE ENOUGH WHEN Scoped episodic/procedural/mission memory with provenance, retention and access policy.

SDK + contract conformance kit

WHY IT MATTERS Necessary for extensibility.

CREDIBLE ENOUGH WHEN Generated client bindings, simulator, manifest validator, conformance tests and reference Specialists.

Resilience / DR / capacity governance

WHY IT MATTERS Required for serious deployment.

CREDIBLE ENOUGH WHEN Defined SLOs, load testing, failure injection, backup, restore and resource budgets.

P2 / Later or deliberately constrained

Marketplace/ecosystem

WHY IT MATTERS Premature before contracts stabilize.

CREDIBLE ENOUGH WHEN Only after manifest/qualification/security contracts are stable.

Continuous autonomous Scout promotion

WHY IT MATTERS High assurance risk.

CREDIBLE ENOUGH WHEN Scout may recommend; promotion requires evidence, testing and governed authorization.

Broad autonomous self-modification

WHY IT MATTERS Conflicts with qualification model if introduced too early.

CREDIBLE ENOUGH WHEN Keep out of core until requalification and rollback are proven.

Full autonomous consequential command

WHY IT MATTERS Not required to demonstrate IQAI's value.

CREDIBLE ENOUGH WHEN Human-authorized action should remain default for high-consequence operations.

Expansive proprietary 3D/robotics engines

WHY IT MATTERS Distracts from core semantics.

CREDIBLE ENOUGH WHEN Integrate specialist engines through contracts instead.

Where IQAI could realistically surpass

IQAI should not target Palantir's broadest horizontal-platform strengths. That is a costly catch-up contest against Ontology, security, connectors, applications, deployment and an established ecosystem simultaneously. Its realistic paths to superiority are narrower, assurance-heavy domains where stronger architectural guarantees matter more than breadth.

Qualified heterogeneous cognition

A first plausible advantage is qualified heterogeneous cognition: a platform in which a local Mission Brain, independently sourced Review Brain, multimodal Perception model, deterministic geospatial engine and domain scientific model can be changed independently, yet each is admitted to a role only after testing against institution-owned competencies. Palantir clearly supports heterogeneous models, but IQAI could exceed current public Palantir architecture by turning cognitive role eligibility into a cross-model, cross-runtime, expiring authorization state. [41]

Forensic epistemic replay

A second is forensic epistemic replay. Conventional observability can answer “what service/model ran?” IQAI should aim to answer “what did the institution know, what was merely reported, what did it infer, which evidence supported each conclusion, which alternative world was active, who challenged it, why was a recommendation authorized, and what happened afterward?” Palantir already has strong lineage, traces and audit, so IQAI would need to prove a substantially richer reason-for-belief layer rather than simply duplicating logging. [53]

Portable institutional competence

A third is portable institutional competence across model generations and vendors. The test would be straightforward: replace the Mission Brain's underlying model with a technologically unrelated local model and demonstrate that the institution's knowledge, procedures, tools, assessments, evidence and mission history remain intact; the new model has no authority until it passes qualification; and rollback does not alter institutional state. Palantir supports model replacement; IQAI needs to demonstrate stronger replacement governance. [54]

Cross-vendor sovereign cognitive federation

A fourth is cross-vendor sovereign cognitive federation. Palantir Peer Manager is a formidable benchmark but operates between Palantir enrollments. IQAI could differentiate if a sovereign IQAI node can federate bounded evidence, world-state assertions or qualified capabilities with another organization while both organizations retain independent infrastructure, identity roots, policies, stores and model choices. [55]

Scientific, deterministic, AI and human authority

A fifth is scientific/deterministic/AI/human authority federation. The win condition is not having Python functions. It is demonstrating that the system knows which conclusions require a certified deterministic calculation, which may be AI-synthesized, which require a qualified specialist, which are hypotheses, and when disagreement forces escalation.


Those are credible “surpass” targets. Connector count, generic dashboards, basic RAG, basic chatbots, generic agents, generic workflow builders, model hosting and conventional GIS are not.

FIGURE 05 / IQAI ARCHITECTURE + BUILD DIRECTION

IQAI — INTELLIGENCE OPERATING ARCHITECTURE

How inputs become governed, evidence-linked intelligence, authorized action and institutional learning

05 / IQAI ARCHITECTURE



TARGET ARCHITECTURE · Built / Partial / Architected / Proposed

05 SEP 2026

Target architecture with mixed maturity. Qualification gates dispatch; consequential actions require human authorization. Academy feedback denotes governed learning and requalification, not automatic unreviewed training.

Source: IQAI-Intelligence-Operating-Architecture-Readable.pptx | Supplied diagram and interpretation notes, 05 September 2026.

05 V3 changes and Palantir interoperability

Keep the topology; harden the contracts

The overall V3 topology should not be replaced. The supplied Evidence Addendum reached the right conclusion: the new evidence strengthens and clarifies V3 more than it changes its topology. The necessary next step is V3.1 constitutional hardening.

Make qualification eligibility a primitive, not a field.

qualification_ref in a CapabilityManifest is not sufficient. Introduce an explicit QualificationGrant or equivalent containing subject binding, role/domain/mission scope, world states, test-suite version, knowledge version, tool versions, policy version, evaluator/authority, issuance, expiry, limitations and revocation/requalification triggers. Routing should consume this object directly. This is the single most important change if IQAI wants the Model–Brain Role distinction to become technology rather than terminology.

Add an Epistemic Promotion state machine.

V3 has strong classes but needs rules for how something moves between them. An Observation should not become an Accepted Baseline Assertion merely because a Brain summarizes it. Define candidate → corroborated → accepted → superseded/revoked states, promotion authority, evidence thresholds and complete provenance. Keep hypothesis/reconstruction/forecast branches non-promotable except through explicit governance.

Freeze WorldStateRef into every relevant envelope.

Evidence results, Specialist outputs, Brain outputs, recommendations and actions should all identify the world they apply to. Otherwise the elegant BASELINE/RECONSTRUCTION/HYPOTHESIS distinction will eventually disappear inside application logic.

Specify a first-class Disagreement object.

A challenger output should not be just another chat message. Store who disagreed, assertion/evidence references, disagreement type, confidence, unresolved questions, arbitration method, disposition and whether the disagreement blocks action. This turns cognitive diversity into inspectable institutional process.

Use a proven durable-execution substrate under Event/Mission Fabric.

V3 should define IQAI mission semantics but deliberately permit Temporal or an equivalent engine as the persistence/replay implementation. Temporal's event-history model is a stronger starting point than constructing another custom workflow database. [\[65\]](#)

Formalize logical-versus-physical Digital World architecture.

The V3 text already reaches the correct conclusion: one logical identity/world system, multiple authoritative physical stores. Make store classes explicit—for example relational/object, graph/index, spatial, time-series, media and evidence blob stores—and prohibit central-world APIs from assuming one backend.

Make sovereign federation a protocol suite.

Separate: trust establishment; node/workload identity; schema negotiation; capability discovery; evidence/assertion exchange; security-label mapping; world-state reconciliation; action delegation; disconnected journal synchronization; conflict resolution; and release synchronization. Palantir Peer Manager proves that mappings, directionality and policy deserve first-class objects. [43]

Keep Academy, Capability Factory and TEVV organizationally and technically distinct.

Academy may identify a competency gap; it should not silently alter production code or promote new models. Capability Factory engineers the change. TEVV independently tests it. Human authority promotes it. This is one of the better safeguards in the Evidence Addendum and should be constitutional.

Add performance and cost to cognitive governance.

Every Brain/Specialist manifest and qualification should carry tested latency/resource envelopes. Selective challenge is essential: invoking every Brain and every Specialist on every mission would create both fan-out latency and cost explosions. Reviewer routing should be risk/evidence/disagreement aware.

Define an open extension ABI at the contract level.

A new Brain, Specialist or deterministic engine should integrate by manifest + typed input/output envelopes + workload identity + qualification + telemetry—not by editing AB. A conformance harness should prove this assertion. Passing that test with technologies from several vendors would be a powerful demonstration of the V3 thesis.

Optional interoperability without dependence

Interoperability with Palantir should be considered a boundary adapter, not an IQAI architecture dependency.

Foundry Ontology

Create an IQAI–Foundry adapter using Palantir APIs/OSDK for authorized object reads or controlled Action requests. Keep IQAI's canonical references, evidence semantics and Brain architecture outside Foundry. Palantir provides both platform APIs and generated Ontology SDKs. [66]

Gotham / maps

Consume or expose authorized layers through OGC-compatible interfaces such as WFS where operationally appropriate; Gaia supports WFS exposure. No IQAI cognitive IP needs to cross the interface. [67]

AIP tools/agents

Expose a bounded IQAI Specialist endpoint or high-level approved capability rather than raw Academy data, prompts, qualification logic or internal orchestration. Conversely, treat a Palantir function/agent as an external Specialist subject to IQAI's own manifest, evidence and qualification requirements.

Data exchange

Exchange curated objects, evidence products or conclusions through APIs/streams rather than giving Palantir direct access to IQAI's Knowledge Vault or internal memory.

Apollo

Package IQAI services as standard signed OCI/Helm-style workloads so a customer that already standardizes on Apollo may deploy them through Apollo. IQAI's ReleaseManifest and trust policy must remain deployment-backend neutral. Apollo's ability to manage heterogeneous disconnected environments makes this a sensible optional enterprise integration. [68]

Identity

Federate at OAuth/OIDC/workload-identity boundaries; do not require IQAI to inherit Palantir's internal identity/control plane.

IP boundary

Export stable service contracts, schemas and decisions—not Mission Brain internals, Academy source corpus, assessment bank, challenger policy, specialist internals or the implementation of qualification/routing.


This allows a future customer to say “use IQAI with our Palantir estate” without turning the product into “IQAI requires Palantir.”

FIGURE 06 / IQAI ARCHITECTURE + BUILD DIRECTION

IQAI — ENGINEERING BUILD MAP

From working specialist systems to a unified, governed intelligence platform

06 / BUILD ROADMAP



06 Delivery roadmap and positioning

CATCH UP → DIFFERENTIATE → SURPASS

CATCH UP

Become operationally credible

DELIVERABLES Universal references; working Evidence Fabric; logical Digital World; durable Mission Fabric; Sensor Gateway; Capability/Brain Registry; model/runtime registry; workload identity/ABAC; ActionEnvelope; full mission tracing; signed deployment/rollback; sovereign local node; SDK baseline

EXIT TEST Run a realistic multi-day mission through failure, restart, model replacement, human wait, disconnected operation and rollback without losing state, provenance or security context.

Prove the production foundation

DELIVERABLES Load/chaos testing, DR, schema migration, connector abstraction, spatial provider federation, audit, policy enforcement, cross-node sync prototype

EXIT TEST Independent red-team/architecture review can reproduce mission history and verify permission boundaries.

DIFFERENTIATE

Make constitutional cognition real

DELIVERABLES Mission/Reviewer BrainRoles; BrainInstance Passport; Specialist Passport; disagreement/arbitration; Academy source → knowledge → assessment → qualification pipeline; qualification-driven routing; explicit world-state separation

EXIT TEST Replace Mission Brain and Reviewer with unrelated models without migrating institutional knowledge; neither may run until separately qualified.

Demonstrate epistemic superiority

DELIVERABLES Claim/evidence graph, promotion/revocation, historical reconstruction, hypothesis isolation, recommendation/decision/action provenance

EXIT TEST A reviewer can trace every high-consequence conclusion to evidence, world state, specialist authority and challenged alternatives.

SURPASS

Lead in assurance-critical heterogeneous intelligence

DELIVERABLES Cross-vendor sovereign cognitive federation; portable institutional competency; selective independent challenge; AI + deterministic + scientific + human Specialist ecosystem

EXIT TEST Multiple organizations with different infrastructure/models collaborate while preserving their own trust roots, evidence control and qualification authority.

Turn architecture into measurable advantage

DELIVERABLES Published qualification benchmarks, replacement/requalification tests, disagreement detection metrics, evidence-grounding rates, mission recovery metrics

EXIT TEST IQAI demonstrates materially lower unsupported-conclusion/action error rates or faster safe model replacement than conventional agent/platform architectures.

The sequencing matters. Do not begin with “surpass.” Without the catch-up layer, the differentiators will sit on unreliable infrastructure. Without the differentiators, IQAI becomes a smaller attempt to reproduce Palantir.

Five comparisons for the private IQAI website

Core architecture

PALANTIR Mature operational platform built around data, Ontology, applications, AI and deployment. [6]

IQAI V3 Designed around a governed Digital World, evidence fabric and federated intelligence constitution.

AI models

PALANTIR Supports multiple commercial, external and self-hosted models across AIP. [69]

IQAI V3 Designed to separate the model from its BrainRole, institutional knowledge, tools, permissions and qualification.

Intelligence governance

PALANTIR Strong platform security, model evaluation, governed Actions and human approvals. [70]

IQAI V3 Designed to add independent Challenger roles, specialist authority and role-specific qualification/requalification.

Knowledge and evidence

PALANTIR Strong Ontology, lineage, audit and governed enterprise context. [71]

IQAI V3 Designed to distinguish evidence, claims, accepted baseline state, hypotheses and recommendations while retaining institution-owned competency outside models.

Strategic focus

PALANTIR Broad enterprise and government operational software platform.

IQAI V3 Focused architectural opportunity: sovereign, evidence-grounded, replaceable and independently qualified heterogeneous intelligence for assurance-critical missions.

The wording “designed to” is important at IQAI’s current maturity. Replacing it with “IQAI provides” before the P0/P1 capabilities are demonstrated would overstate the evidence in the supplied documents.

Investor explanation

Palantir is a mature operational software platform that connects enterprise and mission data to an Ontology, applications, AI models, governed actions and deployment infrastructure. IQAI is not trying to win by rebuilding all of that. Its proposed advantage is a narrower intelligence architecture designed for environments where evidence, sovereignty, independent review and controlled authority matter deeply. In IQAI, a model is not the institutional brain: models are replaceable technologies assigned to governed BrainRoles. Knowledge, memory, tools, permissions, evidence and qualification remain outside the model.

A Mission Brain can therefore use one model, an independent Reviewer another, and a perception specialist another, while each remains separately qualified for its job. IQAI also distinguishes accepted world state from reconstruction, hypothesis and forecast. Palantir is materially ahead in operational maturity today. IQAI’s opportunity is to become stronger at portable, evidence-grounded, independently challenged and continuously requalified institutional intelligence, rather than attempting to become another general-purpose Palantir. [72]

Final architecture recommendation

Freeze the V3 topology; do not freeze the current implementation diagram. The architecture's most valuable intellectual property is not FUNNEL, Digital World, AB or Academy in isolation. It is the set of constitutional invariants between them:

ModelArtifact ≠ BrainRole ≠ BrainInstance ≠ SpecialistCapability ≠ Tool ≠ Knowledge ≠ Memory.

Artifact ≠ Observation ≠ Claim ≠ Evidence ≠ Accepted Baseline ≠ Hypothesis ≠ Recommendation ≠ Decision ≠ Action.

Available capability ≠ authorized capability ≠ qualified capability.

Mission reasoning ≠ independent challenge ≠ specialist authority ≠ human authorization.

Logical world identity ≠ one physical database.

Sovereignty ≠ merely running an LLM locally.

Continuous learning ≠ uncontrolled production self-modification.

Those invariants are where IQAI can still create a coherent architecture that is meaningfully different from Palantir. Palantir has already disproved most simpler claims. It has a mature governed Ontology, heterogeneous models, agents, tools, evaluations, scenario mechanisms, human-approved actions, secure federation and sovereign deployment technology. [73]

Accordingly, the recommended strategy is:

CATCH UP on reliable world state, evidence infrastructure, durable orchestration, security, sovereign deployment, federation, observability and developer tooling.

DIFFERENTIATE through epistemic contracts, explicit world-state separation, qualified BrainRoles, independent Challenger intelligence, typed Specialist authority, institution-owned competency and qualification-driven routing.

SURPASS only where those combinations produce measurable operational outcomes: safer model replacement, better detection of unsupported conclusions, stronger forensic replay, stronger sovereign collaboration and more dependable use of heterogeneous AI/scientific/human intelligence.

That is a much narrower claim than “IQAI is more advanced than Palantir.” It is also a substantially more credible architecture strategy.

07 External architecture benchmarks

The external benchmark set should stay intentionally small. Each system below teaches a specific lesson that Palantir alone does not fully cover.

Esri ArcGIS

IQAI AREA GIS/GEOINT, imagery, 3D, spatiotemporal analytics

WHY IT MATTERS ArcGIS has mature multidimensional raster analysis, time-aware voxel representations and real-time stream services. [56]

WHAT TO LEARN Treat geometry, raster, imagery, terrain, 3D/voxel and streaming geospatial data as distinct optimized physical representations beneath common spatial semantics. Avoid a single “geospatial database.” Adopt open service boundaries and provider abstraction.

OGC STAC

IQAI AREA EO/imagery catalogue federation

WHY IT MATTERS STAC standardizes metadata/querying for satellite, aircraft, drone, hyperspectral, SAR, video, point cloud, lidar, DEM and other spatiotemporal assets, using a minimal extensible core. [57]

WHAT TO LEARN Do not put giant EO assets into Digital World. Keep durable references/catalogue metadata in IQAI and federate assets through STAC-compatible catalogues—exactly consistent with V3's own correction.

Anduril Lattice

IQAI AREA Mission systems, sensors, tactical edge

WHY IT MATTERS Lattice is explicitly built around sensor fusion, mission autonomy and an edge mesh; Mission Autonomy translates intent into tasks and performs distributed sensing/resource management, while Lattice Mesh prioritizes data in degraded environments. [58]

WHAT TO LEARN Learn edge-first processing, bandwidth-aware prioritization, sensor-to-track fusion, local operation and intent/task separation. IQAI's Sensor Gateway and Mission Fabric should be designed for degraded connectivity from the beginning, not as cloud software later made offline.

EXTERNAL BENCHMARKS / CONTINUED

Databricks Unity Catalog + MLflow

IQAI AREA Data/model governance and lifecycle

WHY IT MATTERS Unity Catalog governs models with centralized access control, audit, lineage and discovery; model versions retain training/evaluation lineage. [59]

WHAT TO LEARN Reuse model-registry concepts rather than invent them. Extend them upward: ModelVersion → BrainInstance → qualification scope. Store every qualification against immutable model/runtime/data/knowledge/configuration references.

Temporal

IQAI AREA Durable Mission/Event Fabric

WHY IT MATTERS Temporal persists workflow state, supports recovery/replay and uses event-sourced workflow histories; its architecture separates deterministic workflow logic from side-effecting Activities. [60]

WHAT TO LEARN This is a stronger benchmark than Palantir for IQAI's core durable-orchestration primitive. Adopt the concepts of persisted execution history, deterministic replay, idempotent Activities, durable timers/human waits and version-safe long-running missions rather than writing bespoke retry logic.

Microsoft Agent Framework

IQAI AREA Multi-agent composition

WHY IT MATTERS Current Agent Framework supports sequential, concurrent, handoff, group-chat and manager-coordinated multi-agent workflows plus human-in-the-loop approvals; workflows support checkpointing and explicit graph control. [61]

WHAT TO LEARN Treat multi-agent orchestration mechanics as commodity. IQAI's value should sit above them in role contracts, independence, authority, qualification and evidence/world-state rules.

EXTERNAL BENCHMARKS / CONTINUED

NIST AI RMF / AIRC / TEVV-Athlon

IQAI AREA Academy, qualification, assurance

WHY IT MATTERS NIST's AIRC explicitly supports AI testing, evaluation, verification and validation; the August 2026 TEVV-Athlon draft targets an extensible assessment framework across statistical ML, LLMs, multimodal and agentic systems. [62]

WHAT TO LEARN Keep Academy competency testing separate from independent TEVV, as the Evidence Addendum already recommends. Map Brain/Specialist Passport evidence to recognized lifecycle risk and measurement concepts rather than creating a self-referential IQAI certification regime.

SPIFFE/SPIRE + NIST Zero Trust

IQAI AREA Workload identity and sovereign federation

WHY IT MATTERS SPIRE assigns workload identities and supports federation between trust domains; NIST SP 800-207 remains the foundational zero-trust architecture reference. [63]

WHAT TO LEARN Give every Brain, Specialist, service and sovereign node cryptographic workload identity independent of network location. Model cross-organization federation as explicit trust-domain federation rather than VPN-level “trusted networks.”

Snowflake Secure Sharing / Clean Rooms

IQAI AREA Cross-organization data federation

WHY IT MATTERS Snowflake can expose governed shared data without conventional copying and provides multi-collaborator clean-room patterns where participants control approved queries/templates. [64]

WHAT TO LEARN Study policy-preserving collaboration without surrendering underlying data. An IQAI federation should often exchange approved assertions/results rather than replicate raw institutional evidence.

A major design pattern emerges across these external systems: the strongest architectures separate stable semantic contracts from optimized specialist engines. Esri does not reduce all geospatial data to one representation; Temporal does not make business logic own persistence/retry plumbing; MLflow does not put model lifecycle inside model weights; SPIFFE does not make application credentials the source of identity. That same separation principle should govern IQAI.

08 Sources and reference links

Reference numbers and destination URLs are retained from the supplied report. Click a citation in the report to jump here; click a source title to open the original page.

[1] [6] [9] [13] [14] [22] [33] [34] [39] [72] [73]

[Overview](#) • [Ontology](#) • [Palantir](#)

[2] [32]

[Overview and getting started](#) • [Overview](#) • [Palantir](#)

[3] [8] [21] [41] [54] [69]

[Bring your own model](#) • [Bring your own model to AIP](#) • [Palantir](#)

[4] [19] [23] [42]

[AIP Evals](#) • [Create an evaluation suite](#) • [Palantir](#)

[5]

[2025](#) • [March 2025](#) • [Palantir](#)

[7] [45]

[Overview](#) • [Dev toolchain](#) • [Palantir](#)

[10] [11] [35]

[Overview](#) • [Data integration](#) • [Palantir](#)

[12] [48] [53] [71]

[Data Lineage](#) • [Branching data lineage](#) • [Palantir](#)

[15]

[Geospatial and geotemporal](#) • [Geotemporal series](#) • [Integrate geotemporal series with the Ontology](#) • [Palantir](#)

[16] [52]

[Global Branching](#) • [Overview](#) • [Palantir](#)

[17]

[Core concepts](#) • [Streams](#) • [Palantir](#)

[18]

[Object Explorer](#) • [Overview](#) • [Palantir](#)

[20] [50]

[Functions](#) • [Overview](#) • [Palantir](#)

[24]

[Object edits and materializations](#) • [Overview](#) • [Palantir](#)

[25]

[Automate](#) • [Overview](#) • [Palantir](#)

[26]

[Object permissioning](#) • [Overview](#) • [Palantir](#)

[27] [37] [68]

[Introduction](#) • [Palantir](#)

[28]

[Peer Manager](#) • [Configure file peering](#) • [Palantir](#)

[29] [44]

[Welcome • Palantir](#)

[30] [46]

[Ontology and AIP observability • Overview • Palantir](#)

[31] [66]

[Software development kit \(SDK\) • API Reference • Palantir](#)

[36] [40]

[Object permissioning • Object security policies • Palantir](#)

[38] [43] [55]

[Peer Manager • Peer objects and links between enrollments • Ontology peering • Palantir](#)

[47]

[Bring your own model • Models backed by compute modules • Palantir](#)

[49] [61]

[Workflow orchestrations in Agent Framework | Microsoft Learn](#)

[51] [59]

[Manage model lifecycle in Unity Catalog | Databricks on AWS](#)

[56]

[Multidimensional raster data—ArcGIS Pro | Documentation](#)

[57]

[SpatioTemporal Asset Catalog \(STAC\) - Open Geospatial Consortium](#)

[58]

[Lattice Mesh | Anduril](#)

[60] [65]

[Durable Execution Solutions | Temporal](#)

[62]

[NIST AI Resource Center - AIRC](#)

[63]

[SPIRE Server Configuration Reference | SPIFFE](#)

[64]

[Sobre o Secure Data Sharing | Snowflake Documentation](#)

[67]

[Map as a WFS Source • API Reference • Palantir](#)

[70]

[AIP Evals • Overview • Palantir](#)
